Amendments to the Specification:

Please replace the paragraph at p. 2, line 27 with the following:

If there is any large number of monitored targets 103, the agents return an enormous amount of information to management service 113 and most of the returned information ends up in management repository 115. Consequently, management repository 115 quickly fills up. To reduce the amount of space required by the returned information in repository 115, management service ~~103~~ 113 periodically aggregates the older returned information to reduce its size. To aggregate the information, management service 103 *rolls up* the older returned information to produce rolled up information which is much smaller than the information it was made from and then replaces the older information with the rolled up information. Thus, in FIG. 1, management repository 115 includes current non-rolled up information 117 and less current rolled up information 119.

Please replace the paragraph at p. 3, line 4 with the following:

Since the rolled up information is on the one hand historical but on the other hand needs to be easily accessible to central console 121, management service 103 uses lossy techniques to do the roll up. FIG. 2 gives a simplified example. The events being monitored in the example are hits on Web pages. For purposes of the example, management repository ~~115~~ 215 is taken to include a page hit table ~~101~~ 201 which has entries that record accesses by users on the World Wide Web to Web pages provided by monitored targets 103. There is a hit entry ~~109~~ 209 for each hit on a page in a monitored target 103. Each entry includes three items of data: the URL (universal resource locator) for the page, a time/date stamp ~~105~~ 205 which indicates when the hit occurred, and the source Internet address ~~107~~ 207 of the entity that made the hit.

Please replace the paragraph at p. 3, line 14 with the following:

Clearly, page hit table ~~101~~ 201 will grow very rapidly. Management service consequently periodically rolls up table ~~101~~ 201 to produce a page hit roll up table ~~111~~ 211 for a period *X*. Roll up table ~~111~~ 211 contains only two columns: one, ~~113~~213, for page URLs and one, ~~115~~215, that indicates the number of hits received on the page during the period *X*. There is only one entry for each of the page URLs in table ~~111~~211,

and the value of field ~~115~~ 215 for the entry is the number of hits experienced by the page in the period *X*. As will be immediately apparent from the foregoing, management service 113 makes table ~~111~~ 211 from table ~~101~~ 201 by making a single entry ~~117~~ 217 in table ~~111~~211 for each of the URLs that is present in table ~~101~~201, counting the number of entries for each URL in table ~~101~~ 201 for the period *X*, and placing the result of the count in no. of hits field ~~115~~215. As will also be immediately apparent, table ~~111~~ 211 is far, far smaller than table ~~101~~201. In the following, tables like page hit roll up table ~~111~~ 211 will be called *aggregation* tables and their entries *aggregated entries*, since each entry in table ~~111~~ 211 may aggregate information from many entries in table ~~101~~201. Further, values such as number of hits ~~115~~ 215 which are made by combining a set of values such that the individual values in the set are lost will be termed herein *metric* values. Other examples of such metric values are averages, maxima, minima, modes, and medians. The meaning of a metric value of course depends on the kind of event. For example, if the event indicates that a condition to which the DBA must respond has arisen, the metric value may indicate the time between the time at which the event occurred and the time at which the DBA responded, and the aggregated metric value may be the average response time.

Please replace the paragraph at p. 3, line 34 with the following:

Of course, table ~~111~~ 211 may itself be rolled up. For example, if the period *X* is one day and there is thus a roll up table ~~111~~ 211 for each day, a weekly roll up table may be made from seven daily tables 111. Again, there would be one entry for each URL upon which there was a hit during the week, and no. of hits ~~115~~ 215 would contain the number of hits for the week. The week tables may be rolled up into month tables, the month tables into year tables, and so on. The creation of any roll up entry may be regarded as a *roll up event* at a roll up level *n*, with the entry created by the roll up being a roll up event entry for level *n* and the roll up at level *n*+1 rolling up the roll up event entries for level *n*.

Please replace the paragraph at p. 4, line 8 with the following:

Aggregation tables are challenging to design. The challenge is to reduce the size of the information in the aggregation table as much as possible while reducing the usefulness of

the information contained in the table as little as possible. Table ~~111~~ 211 illustrates the difficulty. In page hit table ~~101~~201, the time at which each hit occurred is recorded in time/date field ~~105~~205; this information is lost in table ~~111~~211; thus, though table ~~111~~ 211 can tell the DBA how many times a page was hit in the period *X*, it cannot tell the DBA anything about the temporal distribution of hits over the period *X*. This pattern information may, however, be exactly what the DBA needs to correctly distribute copies of the page among monitored targets 103. What is needed if rollup table ~~111~~ 211 is to provide useful information about the temporal distribution of the hits is a way of representing time/date information 105 in aggregated page hit entry ~~117~~ 217 for the individual hit entries ~~109~~ 209 that have been aggregated into entry ~~117~~217. In more general terms, the problem is this: how to incorporate information that consists of sets of values into aggregation tables. What is needed, and what is provided by the invention disclosed herein, is a technique for doing this.

Please replace the paragraph at p. 6, line 23 with the following:

Modifying page hit rollup table ~~111~~ 211 to specify the times that the hits occurred: FIG. 3

Please replace the paragraph at p. 6, line 24 with the following:

FIG. 3 shows a version 301 of page hit rollup table ~~111~~ 211 which has been modified to include a representation of the time/date information contained in field ~~105~~ 205 of page hit table ~~101~~201. The representation makes up column 303 of table 301. The data in fields belonging to column 303 represents the time/date information for all of the records from table ~~101~~ 201 that are aggregated into an entry in table 301 as a set of the times at which the hits occurred. Thus, for the record that aggregates the hits on page URL A, the set value is the set $\{x\}$; for the record that aggregates the hits on page URL B, the set value is the set $\{y\}$. In both cases, the time at which each hit on the record's URL occurred is a member of the set represented by the set value. The set value is made when the records aggregated into the entry in page hit roll up table 301 are aggregated; each time one of the records being aggregated is read, its time/date stamp ~~105~~ 205 is made into a member of the set represented by field 303 in the aggregated record. Of course, set

values in a rollup record may be further rolled up. For example, the set of times representing hits during a day may be rolled up into a set of times representing hits during a week. Further, the granularity of the set may be reduced in such a rollup. For instance, the weekly rollup may include a set value in which members of the set specify the number of hits per hour for each hour of the week.

Please replace the paragraph at p. 8, line 17 with the following:

The time/date information from page hit table ~~101~~ 201 is particularly easy to represent in set data items because the information for a given page URL ~~103~~ 203 forms a monotonically increasing set of values. However, other kinds of values can be represented in set values as well. For example, the comma list can be used with values that do not increase monotonically or for sets where the values are tuples instead of single numbers. For example, a weekly rollup might represent the hit times with a comma list like this: *. . .; day of week, second in day; . . .* where the entries in the list are separated by semicolons and the tuple in the entry specifies the ~~the~~ day of the week and the second the hit occurred on that day. Comma lists with tuples can similarly be used to represent sets of coordinates in two and three dimensions. Where nested tables are possible, the nested tables could be used in place of the comma lists. Where a set never has two members with the same value and the number of possible values is finite, the set may be mapped to a bit set data item as described above for time values. The bit set data item can of course have more than one dimension; thus, a bit set data item might be used in the weekly rollup to represent a histogram of the hits that occurred during the hours making up the week.

Please replace the paragraph at p. 11, line 25 with the following:

Entry 517(k) is a monthly roll up entry in the alert history table for the month of May. It is made beginning at midnight on ~~June~~ May 1. Entry 517 has the same form as entries 501, except that `RollupInterval` 513 indicates that the roll up window is one month and set value field 519 contains the times from the daily rollup entries that were rolled up into the monthly roll up entry.